

Getting Started with Amazon EC2

David J. Malan

malan@post.harvard.edu

www.cs50.net

Self-service, Prorated Super Computing Fun!

- “The New York Times has decided to make all the public domain articles from 1851-1922 available free of charge. These articles are all in the form of images scanned from the original paper. In fact from 1851-1980, all 11 million articles are available as images in PDF format. To generate a PDF version of the article takes quite a bit of work — each article is actually composed of numerous smaller TIFF images that need to be scaled and glued together in a coherent fashion. . . . I then began some rough calculations and determined that if I used only four machines, it could take some time to generate all 11 million article PDFs. But thanks to the swell people at Amazon, I got access to a few more machines and churned through all 11 million articles in just under 24 hours using 100 EC2 instances (In fact, it work so well that we ran it twice, since after we were done we noticed an error in the PDFs.)”

Virtualization



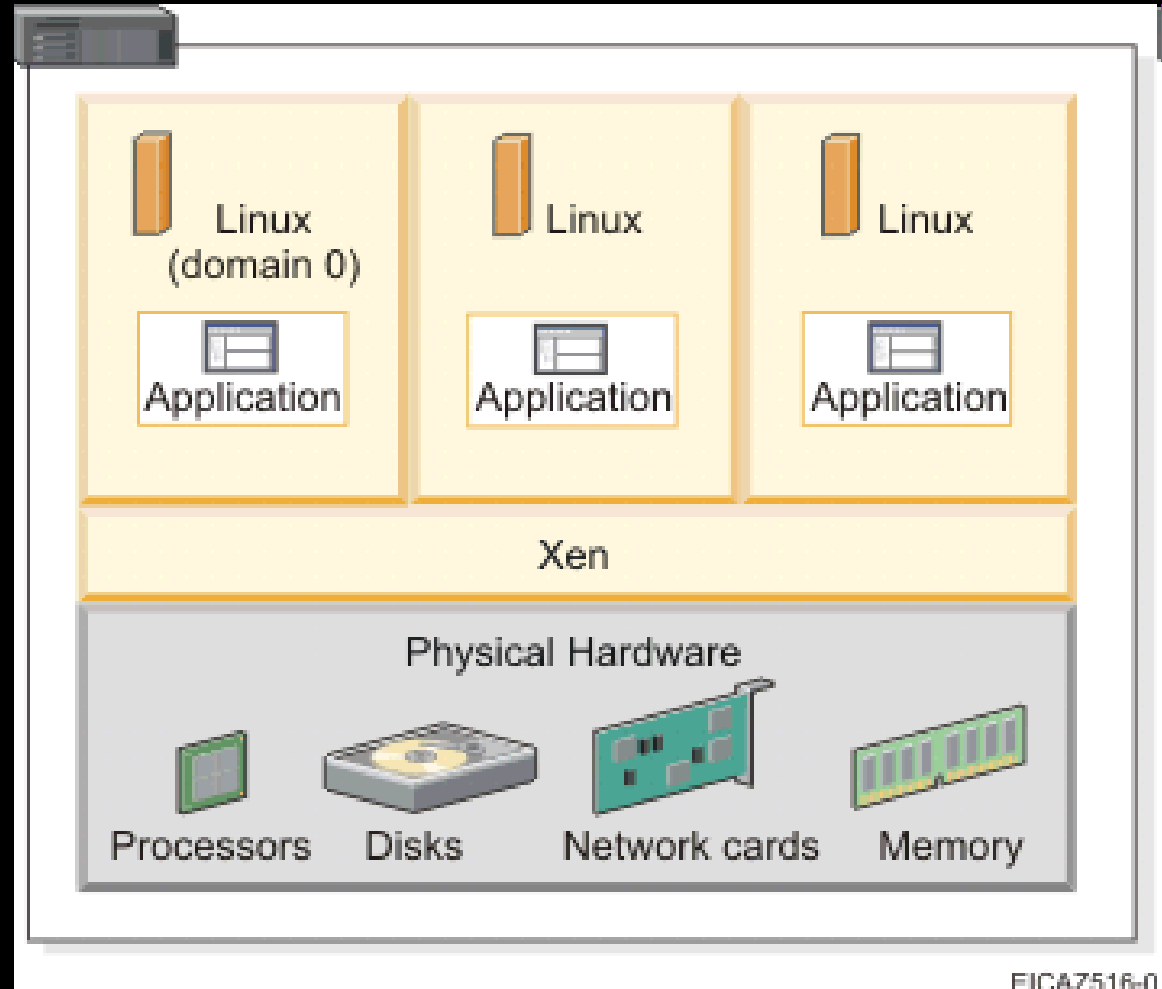
Virtualization

- **Xen**
- VMware
 - Workstation, Fusion, Server, ESXi, vSphere
- Parallels
 - Workstation, Server for Mac, Virtuozzo
- Microsoft
 - Hyper-V, Virtual PC
- Sun
 - VirtualBox
- QEMU
- FreeVPS
- Virtual Iron
- ...

Virtualization

- Full Virtualization
- Paravirtualization
- OS-Level Virtualization

Xen



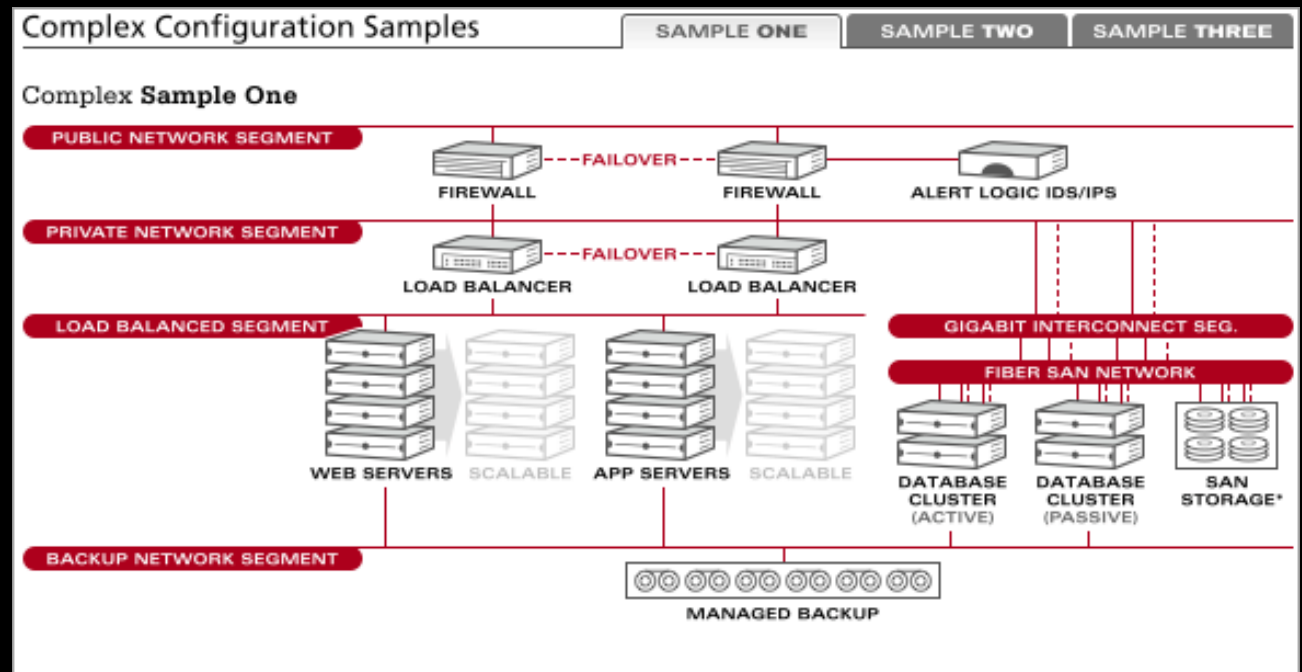
Virtual Private Servers (VPSes)

- Linode
- ServInt
- TekTonic
- VPSLAND
- ...

SuperVPS	SuperVPS 125	SuperVPS 2x
75 GB Disk Space 3 TB Monthly Transfer 4 GB Guaranteed (8 GB Burst RAM) 1 CPU Core Guaranteed (Burst to 4 Cores) Hardware RAID 10 CentOS 5 Operating System 4 IP Addresses Unlimited Domains and User Accounts FREE Virtuozzo Power Panel FREE nightly backups FREE cPanel or Plesk 9 Available [details] No set-up fee \$199 monthly Features, Upgrades and Options	125 GB Disk Space 3.5 TB Monthly Transfer 4 GB Guaranteed (8 GB Burst RAM) 1 CPU Core Guaranteed (Burst to 4 Cores) Hardware RAID 10 CentOS 5 Operating System 4 IP Addresses Unlimited Domains and User Accounts FREE Virtuozzo Power Panel FREE nightly backups FREE cPanel or Plesk 9 Available [details] No set-up fee \$229 monthly Features, Upgrades and Options	150 GB Disk Space 6 TB Monthly Transfer 8 GB Guaranteed (16 GB Burst RAM) 2 CPU Cores Guaranteed (Burst to 8 Cores) Hardware RAID 10 CentOS 5 Operating System 4 IP Addresses Unlimited Domains and User Accounts FREE Virtuozzo Power Panel FREE nightly backups FREE cPanel or Plesk 9 Available [details] No set-up fee \$349 monthly Features, Upgrades and Options
ORDER	ORDER	ORDER

Managed Colocation

- Rackspace
- ...



Clouds


- Amazon Elastic Compute Cloud (Amazon EC2)
- Google App Engine
- Microsoft Azure Services Platform
- ...

Amazon Web Services (AWS)

- **Amazon Elastic Compute Cloud (Amazon EC2)**
- **Amazon Simple Storage Service (Amazon S3)**
- Amazon SimpleDB
- Amazon CloudFront
- Amazon Simple Queue Service (Amazon SQS)
- Amazon Elastic MapReduce



cloud.cs50.net



Jargon

- AMI
- Instance
- EBS Volume
- ...

AMIs

- Debian
- Fedora
- Gentoo Linux
- Red Hat Enterprise Linux
- OpenSolaris
- openSUSE Linux
- Oracle Enterprise Linux
- Ubuntu Linux
- Windows Server 2003
- ...

Standard Instances

- **Small Instance (m1.small)**
 - 1.7 GB memory
 - 1 EC2 Compute Unit (1 virtual core with 1 EC2 Compute Unit)
 - 160 GB instance storage (150 GB plus 10 GB root partition)
 - 32-bit platform
- **Large Instance (m1.large)**
 - 7.5 GB memory
 - 4 EC2 Compute Units (2 virtual cores with 2 EC2 Compute Units each)
 - 850 GB instance storage (2×420 GB plus 10 GB root partition)
 - 64-bit platform
- **Extra Large Instance (m1.xlarge)**
 - 15 GB memory
 - 8 EC2 Compute Units (4 virtual cores with 2 EC2 Compute Units each)
 - 1,690 GB instance storage (4×420 GB plus 10 GB root partition)
 - 64-bit platform

High-CPU Instances

- **High-CPU Medium Instance (c1.medium)**
 - 1.7 GB of memory
 - 5 EC2 Compute Units (2 virtual cores with 2.5 EC2 Compute Units each)
 - 350 GB of instance storage
 - 32-bit platform
- **High-CPU Extra Large Instance (c1.xlarge)**
 - 7 GB of memory
 - 20 EC2 Compute Units (8 virtual cores with 2.5 EC2 Compute Units each)
 - 1690 GB of instance storage
 - 64-bit platform

Regions + Availability Zones

- **United States**
 - us-east-1a
 - us-east-1b
 - us-east-1c
- **Europe**
 - eu-west-1a
 - eu-west-1b

On-Demand Instances

United States	Europe	
Standard On-Demand Instances	Linux/UNIX Usage	Windows Usage
Small (Default)	\$0.10 per hour	\$0.125 per hour
Large	\$0.40 per hour	\$0.50 per hour
Extra Large	\$0.80 per hour	\$1.00 per hour
High CPU On-Demand Instances	Linux/UNIX Usage	Windows Usage
Medium	\$0.20 per hour	\$0.30 per hour
Extra Large	\$0.80 per hour	\$1.20 per hour

Reserved Instances

United States	Europe		
Linux/UNIX			
One-time Fee			
Standard Reserved Instances	1 yr Term	3 yr Term	Usage
Small (Default)	\$325	\$500	\$0.03 per hour
Large	\$1300	\$2000	\$0.12 per hour
Extra Large	\$2600	\$4000	\$0.24 per hour
High CPU Reserved Instances	1 yr Term	3 yr Term	Usage
Medium	\$650	\$1000	\$0.06 per hour
Extra Large	\$2600	\$4000	\$0.24 per hour

Data Transfer

Data Transfer In

All Data Transfer	\$0.10 per GB
-------------------	---------------

Data Transfer Out

First 10 TB per Month	\$0.17 per GB
-----------------------	---------------

Next 40 TB per Month	\$0.13 per GB
----------------------	---------------

Next 100TB per Month	\$0.11 per GB
----------------------	---------------

Over 150 TB per Month	\$0.10 per GB
-----------------------	---------------

Data Transfer

- **Availability Zone Data Transfer**
 - \$0.00 per GB
- **Regional Data Transfer**
 - \$0.01 per GB in/out
- **Public and Elastic IP Data Transfer**
 - \$0.01 per GB in/out
- **Elastic IP Addresses**
 - No cost while in use
 - \$0.01 per non-attached Elastic IP address per complete hour
 - \$0.00 per Elastic IP address remap – first 100 remaps / month
 - \$0.10 per Elastic IP address remap – additional remap / month over 100

EBS Volumes + Snapshots

United States

Europe

Amazon EBS Volumes

- \$0.10 per GB-month of provisioned storage
- \$0.10 per 1 million I/O requests

Amazon EBS Snapshots to Amazon S3 (priced the same as Amazon S3)

- \$0.15 per GB-month of data stored
- \$0.01 per 1,000 PUT requests (when saving a snapshot)
- \$0.01 per 10,000 GET requests (when loading a snapshot)

Elastic IP Addresses

- No cost while in use
- \$0.01 per non-attached Elastic IP address per complete hour
- \$0.00 per Elastic IP address remap – first 100 remaps / month
- \$0.10 per Elastic IP address remap – additional remap / month over 100

Cost Less than \$15 per Student

	CPU	Disk	I/O Requests	Bandwidth
Sep	2,275 Hrs	125 GB	45,348	14 GB
Oct	3,425 Hrs	108 GB	93,257,314	191 GB
Nov	5,484 Hrs	199 GB	337,019,916	239 GB
Dec	5,206 Hrs	300 GB	427,639,962	52 GB
Jan	5,208 Hrs	300 GB	1,502,614,186	62 GB

AWS Simple Monthly Calculator

<http://calculator.s3.amazonaws.com/calc5.html>

Concerns

- Time
- No _____
- PEBKAC
- “cloud is laggy”
- “Cloud is wicked laggy”
- Bandwidth Costs
- “Lightning Strike Triggers Amazon EC2 Outage”

CS 50 in a Box



Resources

- **Getting Started Guide**
 - <http://docs.amazonwebservices.com/AWSEC2/latest/GettingStartedGuide/>
- **User Guide**
 - <http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/>
- **Developer Guide**
 - <http://docs.amazonwebservices.com/AWSEC2/latest/DeveloperGuide/>
- ...

AWS in Education

<http://aws.amazon.com/education/>

- Educators
- Researchers
- Students
- Education IT

What's Next?

- Q&A
- Getting Your Hands Dirty
- Getting Your Hands Dirtier

Getting Started with Amazon EC2

David J. Malan

malan@post.harvard.edu